

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
18.01.2006 Bulletin 2006/03

(51) Int Cl.:
H04L 9/32 (2006.01)

(21) Application number: 04405442.7

(22) Date of filing: 12.07.2004

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR
HU IE IT LI LU MC NL PL PT RO SE SI SK TR
Designated Extension States:
AL HR LT LV MK

- Camenisch, Jan
8803 Rüschlikon (CH)
- Schunter, Matthias
8003 Zurich (CH)
- Waidner, Michael
8804 Au (CH)

(71) Applicant: International Business Machines
Corporation
Armonk, N.Y. 10504 (US)

(74) Representative: Toleti, Martin
IBM Research GmbH
Zurich Research Laboratory
Säumerstrasse 4 / Postfach
8803 Rüschlikon (CH)

(72) Inventors:
• Bangerter, Endre-Felix
2514 Ligerz (CH)

(54) **Method, system and computer program product for privacy-protecting integrity attestation of computing platform**

(57) The method for privacy-protecting integrity attestation of a computing platform (P) having a trusted platform module (TPM) comprises the following steps. First, the computing platform (P) receives configuration values (PCR1 ... PCRn). Then, by means of the trusted platform module (TPM) a configuration value (PCRp) is determined which depends on the configuration of the

computing platform (P). In a further step the configuration value (PCRp) is signed by means of the trusted platform module (TPM). Finally, in the event that the configuration value (PCRp) is one of the received configuration values (PCR1 ... PCRn), the computing platform (P) proves to a verifier (V) that it knows the signature (sign(PCRp)) on one of the received configuration values (PCR1 ... PCRn).

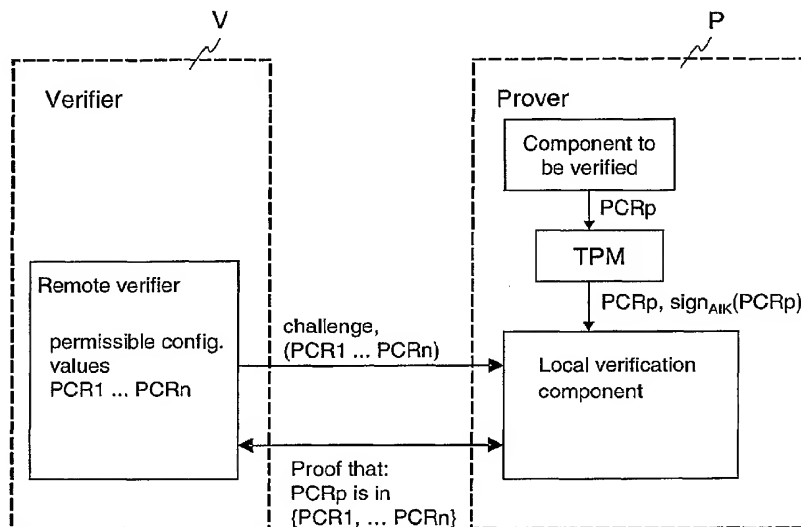


Fig. 3

Description

TECHNICAL FIELD

5 [0001] The present invention relates to a method for privacy-protecting integrity attestation of a computing platform, a computing platform for privacy-protecting integrity attestation, and a network for privacy-protecting communication.

BACKGROUND OF THE INVENTION

10 [0002] Processing critical information relies on the security of the computing platform. Typical security goals are to prevent such critical information from leaking beyond the realm of machines that are trusted by the user or to prevent corrupted machines from impacting the integrity of a computation. External verification of platform integrity enables a machine to verify that another machine meets certain security requirements. This is useful, for example, when a grid server wants to assure that a grid node is untampered before delegating a grid process to it.

15 [0003] In the following example two computers or machines A and B interact over a network. In the example machine A can make certain statements about its own state, e.g. "I am in status access" or "My software ... is in status x", or deliver a hash or checksum of this status ($h(x)$), or certain properties, e.g. "I am running version ... of Linux". Machine A can send these statements to machine B, but why should machine B trust machine A with respect to the correctness of these statements? If machine A is corrupted by a hacker, it could make arbitrary claims about itself.

20 [0004] Therefore, it is desired to implement a proving method with which machine B can verify whether the statements made by machine A are correct. The embodiment of such a proving method is shown in FIG. 1. In the following, machine A is called verified machine or the prover and machine B the verifier. All solutions to this problem assume that there is a piece of hardware, called trusted platform module TPM, which cannot be compromised, and which can make reliable statements about the rest of the system A. Specifically, the industry consortium Trusted Computing Group (TCG) has specified the trusted platform module TPM, which can compute a checksum of the system configuration of machine A, wherein the checksum can be computed for a system configuration in which all or only a part of the software is running on machine A. In a further step the computed checksum is signed, and afterwards send off to the verifier B. The corresponding protocol is shown in FIG. 2.

30 [0005] The Trusted Computing Group is an IT industry consortium which has developed a specification of a small, low-cost commodity hardware module, called trusted platform module (TPM). The TPM can serve as a root of trust in remote (and local) platform verification. The base TCG model of this configuration verification process, called binary attestation, aims at measuring all executed code. Therefore, each measured piece of software stores metrics of a sub-component into the TPM before executing it, wherein the metrics are hash values of the configuration's components. The metrics are bootstrapped by the basic input output system (BIOS) that is trusted by default and that is measuring and storing the boot loader. The chain of trust can then be extended to the operating system components and to the applications and their configuration files. Once the executables are measured into the TPM, the TPM can reliably attest to the metrics of the executed components by signing the metrics with a TPM-protected key. The signed metrics, also called integrity metrics, can then be transmitted to a verifying machine. This verifier machine, or in short verifier, can decide whether to consider the verified machine trustworthy enough to involve it in a subsequent computation. As will be elaborated hereinafter, this straightforward approach of binary attestation lacks privacy. The main reason therefore is that the whole configuration is transmitted.

45 [0006] Hereinafter, the binary attestation and verification is explained. The ability of the TPM reliably to report on the verified platform's computing environment follows from the TPM-enabled measurement and reporting. The description in the following paragraphs focuses on a PC-platform. Nevertheless, the TPM can be implemented in any platform or computing device, e.g. mobile phone, PDA, notebook, etc. The measurement and storage of integrity metrics is started by the BIOS boot block (a special part of the BIOS which is believed to be untampered) measuring itself and storing the measurements in a TPM PCR (platform configuration register) before passing control to the BIOS. In the same way, the BIOS then measures option ROMs and the boot loader and records these measurements in a TPM PCR before passing control to the boot loader. The process continues as the boot loader measures and stores integrity metrics of the operating system (OS) before executing it. The OS in turn measures and stores integrity metrics of additionally loaded OS components before they are executed. If support by the OS is provided, applications can also be measured before being executed. The measurement and reporting processes are depicted in a simplified manner in FIG. 2, in which H represents the cryptographic hash function SHA-1. During initialization, various platform configuration registers PCR_x as well as a configuration log file *log* (stored on the platform) are initialized. This log file *log* keeps track of additional information such as descriptions or file paths of loaded components. Its integrity needs not to be explicitly protected by the TPM. 55 During subsequent measurement of components, this log file *log* is extended, while metrics (hash values) of the executables are stored in the TPM using the *tpm_extend* method replacing the contents of the appropriate platform configuration register PCR_x with the hash of the old contents and the new metrics, wherein metrics of loaded components are reliably stored in the TPM. When a remote verifier B wants to assess the security of the verified platform A, the verifier

B sends a challenge c to the platform A. The platform A uses this challenge c to query with a `tpm_quote` command the TPM for the value of the platform configuration registers PCR. The TPM responds with a signed message $qu = \text{sign}_{\text{AIK}}$

($\overrightarrow{\text{PCR}}, c$) containing the PCR values and the challenge c . The platform A returns this signed quote qu to the challenger

(verifier B) together with information from the log file *log* needed by the verifier B to reconstruct the verified platform's configuration. The verifier B can then decide whether this configuration is acceptable. The key used for signing the quote is an attestation identity key AIK of the TPM. As a TPM may have multiple attestation identity keys, the key or its identifier has to be specified in the `tpm_quote` request. An attestation identity key AIK is bound to a specific TPM. Its public part is certified in an attestation identity key certificate by a privacy-certification authority as belonging to a valid TPM. The verifier of a quote signed with a correctly certified AIK believes that the quote was produced by a valid TPM, more specifically, by the unique TPM owning that AIK. This belief is based on the assumption that the TPM is not easily subject to hardware attacks and that effective revocation mechanisms are in place dealing with compromised keys.

[0007] Note that the above measurement process does not prohibit execution of untrusted code, it only guarantees that the measurement of such code will be securely stored in the TPM. Thus, if malicious code is executed, the integrity of the platform A may be destroyed. However, the presence of an untrusted or simply unknown component will be reflected by the TPM quotes not matching the correct or expected values.

[0008] Further information about the Trusted Computing Group and the trusted platform module can be found in The Trusted Computing Group, Main specification version 1.1b, 2003, which is available from <http://www.trustedcomputing-group.org>.

[0009] The trusted platform module TPM also supports trusted booting, which means that the prover A can go through a sequence of steps. In each step a new component is loaded e.g., first the boot loader, then the operating system, and then an application. The TPM ensures that each step succeeds only if all previous steps succeeded. Via trusted booting one can ensure that the prover A boots into a known, well-defined state, with certain properties.

[0010] A related, theoretically well investigated feature is secure booting. The difference to trusted booting is that a system with secure booting either boots a specific, pre-defined system or does not boot at all, while a system with trusted booting can boot any system, but certain data are accessible only if it boots into a pre-defined system. The details of how a secure boot process can be carried out can be looked up in B. Yee, "Using secure coprocessors", Technical Report CMU-CS-94-149, Carnegie Mellon University School of Computer Science, May 1994.

[0011] The above mentioned TCG specifications define mechanisms for a TPM-enabled platform to reliably report its current hardware and software configuration to a local or remote challenger. This binary attestation, based on measurements of binary executables, firstly requires that the platform builds a chain of trust from the hardware up to the operating system and potentially, including applications by measuring integrity metrics of modules and storing them in the TPM, and secondly requires that the TPM is able to report on these integrity metrics in an authenticated way. A verifier obtaining such authenticated integrity metrics can then match them against the values of a known configuration and decide whether the verified machine meets the security requirements or not. But with that, the privacy of the verified machine is violated, because with this approach the prover needs to reveal its configuration. The verifier gets information about the configuration of the verified machine, e.g. which operating system and which applications are running on the verified machine. This will not be acceptable to consumer since it raises substantial privacy concerns.

SUMMARY OF THE INVENTION

[0012] Therefore, an object of the invention is to provide a method for privacy-protecting attestation of a computing platform, where the computing platform can keep its configuration secret. That is the computing platform does not have to reveal its configuration.

[0013] According to one aspect of the invention, the object is achieved by a method for privacy-protecting integrity attestation of a computing platform with the features of the independent claim 1.

[0014] The method for privacy-protecting integrity attestation of a computing platform having a trusted platform module comprises the following steps. First, the computing platform receives configuration values. Then, by means of the trusted platform module a configuration value is determined which depends on the configuration of the computing platform. In a further step the configuration value is signed by means of the trusted platform module. Finally, in the event that the configuration value is one of the received configuration values, the computing platform proves to a verifier that it knows the signature on one of the received configuration values.

[0015] According to another aspect of the invention, the object is achieved by a computing platform for privacy-protecting integrity attestation.

[0016] The computing platform for privacy-protecting integrity attestation according to the invention is formed such that it is able to receive configuration values and comprises a trusted platform module which in turn is formed such that it is able to determine a configuration value depending on the configuration of the computing platform, and to sign the configuration value. Furthermore, the computing platform is formed such that in the event that the configuration value

is one of the received configuration values it is able to prove to a verifier that it knows the signature on one of the received configuration values.

[0017] According to a further aspect of the invention, the object is achieved by a network for privacy-protecting communication.

[0018] The network for privacy-protecting communication according to the invention comprises a verifier, and a computing platform having a trusted platform module, wherein the computing platform is formed such that it is able to receive configuration values, e.g. from the verifier. The trusted platform module is formed such that it is able to determine a configuration value depending on the configuration of the computing platform, and to sign the configuration value. The computing platform is furthermore formed such that in the event that the configuration value is one of the received configuration values it is able to prove to the verifier that it knows the signature on one of the received configuration values.

[0019] Advantageous further developments of the invention arise from the characteristics indicated in the dependent patent claims.

[0020] Preferably, in the method according to the invention the proof is carried out by a cryptographic proof.

[0021] Advantageously, in the method according to the invention the proof is carried out by a zero-knowledge proof.

With that, it can be ensured that the information the verifier gets from the computing platform does not contain any knowledge about the configuration of the computing platform itself.

[0022] In an embodiment of the method according to the invention the computing platform checks after receiving the configuration values whether configurations having the received configuration values actually may exist, and if this is the case the above described method is further processed. With that, the privacy-protection can be further improved.

[0023] In another embodiment of the method according to the invention the computing platform checks whether the number of received configuration values exceeds a minimum number of configuration values, and if this is the case the above described method is further processed. With that, the privacy-protection can be further improved.

[0024] The present invention also relates to a computer program element comprising program code for performing the above described method when loaded in a digital processor of a computer.

[0025] Finally, the present invention also relates to a computer program product stored on a computer usable medium, comprising computer readable program code for causing a computer to perform the above described method.

[0026] Additional objects and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0027] The invention and its embodiments will be more fully appreciated by reference to the following detailed description of presently preferred but nonetheless illustrative embodiments in accordance with the present invention when taken in conjunction with the accompanying drawings.

[0028] The figures are illustrating:

FIG. 1 a block diagram of the architecture for a system for binary attestation according to the prior art,

FIG. 2 a protocol for the architecture shown in FIG. 1, and

FIG. 3 a block diagram of the architecture for a system for privacy-protecting attestation of the integrity of a computing platform according to the invention.

DETAILED DESCRIPTION OF AN EMBODIMENT

[0029] The description of the FIGs. 1 and 2 is in section "Background of the invention".

[0030] Below is described a way how to implement privacy-friendly attestation where the prover and the verifier ensure that the prover has one of the given configurations without revealing which one. As a consequence, the verifier learns nothing except that the prover has a correct or permissible configuration. This enhances the privacy of the prover substantially.

[0031] In FIG. 3 a block diagram of the architecture for a system for privacy-protecting attestation of the integrity of a computing platform according to the invention is depicted. A computing platform P, in the following also called platform or prover, comprises a trusted platform module TPM that measures the platform P. This measurement results in a platform configuration register value stored in a platform configuration register (PCR). Different PCR values correspond to different configurations of the platform P. To convince the verifier V that the platform P is in a certain configuration, the TPM signs the measured PCR value PCR_p with respect to an attestation identity key AIK which comprises an RSA public/secret key pair. Afterwards, the TPM sends the PCR value PCR_p and its signature sign_{AIK}(PCR_p) to the local verification component, which is a part of the platform P. It receives the PCR value PCR_p and its signature sign_{AIK}(PCR_p) and then runs the protocol below. That is, the platform P is given the PCR value, the RSA signature on the PCR value, and the RSA public key part of the attestation identity key AIK. The platform P wants to produce from this a proof that

it knows an RSA signature with respect to this attestation identity key AIK on, e.g., one out of a list of PCR values. In the following it is described how this can be achieved.

[0032] By using a cryptographic proof, particularly a zero-knowledge proof, the prover P can prove its integrity to the verifier without revealing its configuration. This means, the prover P can proof to the verifier V by means of the cryptographic proof that its configuration is indeed one of the configurations the verifier V has specified as being correct or permissible.

[0033] The protocol described below uses techniques devised by Louis C. Guillou and Jean-Jacques Quisquater, "A "paradoxical" identity-based signature scheme resulting from zero-knowledge" in Shafi Goldwasser, editor, Advances in Cryptology - CRYPTO '88, volume 403 of LNCS, pages 216 - 231, Springer Verlag, 1988, Ronald Cramer, Ivan Damgård, and Berry Schoenmakers "Proofs of partial knowledge and simplified design of witness hiding protocols" in Yvo G. Desmedt, editor, Advances in Cryptology - CRYPTO '94, volume 839, pages 174 - 187, Springer Verlag, 1994, and Alfredo De Santis, Giovanni Di Crescenzo, Giuseppe Persiano, and Moti Yung "On monotone formula closure of SZK" in 35th Annual Symposium on Foundations of Computer Science, pages 454 - 465, Santa Fe, New Mexico, 20-22 November 1994, IEEE. In the following key notions are briefly recalled and some notations used are described. For details reference is made to the above mentioned papers.

[0034] Let θ denote a smooth secret sharing scheme and s a secret that is shared among n participants, each having an element of the set of secret shares $\{s_1, \dots, s_n\}$. A qualified set of secret shares for θ is a set of shares that allows to reconstruct s . The access structure Γ is the set of index sets that contain the indices of secret shares being qualified sets.

[0035] The dual access structure to the access structure r is denoted by Γ^* and defined by:

$$A \in \Gamma^* \Leftrightarrow \bar{A} \notin \Gamma$$

[0036] By definition of a smooth secret sharing scheme θ there exist algorithms as follows:

[0037] A first algorithm is called "is-consistent ($s, \{s_1, \dots, s_n\}, \Gamma$)", whereas the first argument s of the algorithm is a secret, the second argument $\{s_1, \dots, s_n\}$ is a full set of secret shares, and the third argument Γ is an access structure. The algorithm is-consistent ($s, \{s_1, \dots, s_n\}, \Gamma$) returns 1, if all qualified sets in r determine s as the secret. Such a triple ($s, \{s_1, \dots, s_n\}, \Gamma$) is called consistent. Otherwise the algorithm returns {}.

[0038] A second algorithm is called "complete ($s, \{s_1, \dots, s_k\}, \Gamma$)", whereas the first argument s of the algorithm is a secret, the second argument $\{s_1, \dots, s_k\}$ is a set of incomplete shares, i.e., a not qualified set of shares, and the third argument r is an access structure. The algorithm complete ($s, \{s_1, \dots, s_k\}, \Gamma$) returns a set of shares S such that $S \cup \{s_1, \dots, s_k\}$ is consistent, whereas $\{s_1, \dots, s_k\} \cap S = \emptyset$.

[0039] This general technique allows one to prove arbitrary statements, for instance "I know a signature on PCR1 and PCR2 or one on PCR2 and PCR3 or one on PCR4." In practice, however, the statement will often be "I know a signature on one out of the following PCR1, PCR2, PCR3, and PCR4." In this case, the secret sharing scheme θ can be realized quite simply: If the secret s is a secret k -bit string, choose $k-1$ shares s_i as random k -bit strings and then set the share s_k to:

$$s_k = s \oplus (\oplus_{i=1}^{k-1} s_i) = s \oplus s_1 \oplus s_2 \oplus \dots \oplus s_{k-1}$$

wherein \oplus denotes a XOR-operation.

That is, the algorithm is-consistent ($s, \{s_1, \dots, s_n\}, \Gamma$) checks whether the secret s holds the condition:

$$s = \oplus_{i=1}^k s_i$$

[0040] The algorithm complete ($s, \{s_1, \dots, s_{j-1}, s_{j+1}, \dots, s_k\}, \Gamma$) sets the share s_j to:

$$s_j = s \oplus (\oplus_{i=1, i \neq j}^k s_i).$$

[0041] In the following is described how the knowledge of a RSA signature can be proved.

[0042] Let m be the message that is signed, σ the signature on the message m , and (e, n) the public key of the signer,

wherein the signer is for example the TPM and in the context of TCG the exponent $e = 2^{16} - 1$. Let H be the encoding function as defined in PKCS#1 V1.5 (RSA Encryption Standard, <http://www.rsasecurity.com/rsalabs/node.asp?id=2125>). The value σ is a valid signature on the message m if $\sigma^e \equiv H(m) \pmod{n}$ with respect to the public key (e, n) .

[0043] Considering now a zero knowledge protocol that allows the prover P to convince the verifier V that it knows a signature σ on a message m with respect to the public key (e, n) where the verifier V is allowed to learn the message m and the public key (e, n) but not the signature σ .

[0044] It is assumed that the zero knowledge protocols run in practice in non-interactive mode. That is, one can use the Fiat-Shamir heuristic with the protocols. The Fiat-Shamir heuristic is described in Amos Fiat and Adi Shamir "How to prove yourself: Practical solutions to identification and signature problems" in Andrew M. Odlyzko, editor, *Advances in Cryptology - CRYPTO '86*, volume 263 of LNCS, pages 186 - 194, Springer Verlag, 1987. In this mode, the challenge c in the protocol is determined via a hash-function on the first message t in the protocol and the protocol's public inputs.

[0045] One way to achieve the required proof is to apply the protocol by Louis C. Guillou and Jean-Jacques Quisquater, "A "paradoxical" identity-based signature scheme resulting from zero-knowledge" in Shafi Goldwasser, editor, *Advances in Cryptology - CRYPTO '88*, volume 403 of LNCS, pages 216 - 231, Springer Verlag, 1988. It works as follows:

1. The prover P carries out the following steps:

The prover P chooses a random value r , wherein $r \in \mathbb{Z}_N^*$, sets the first message t to:

$$t = r^e \pmod{N},$$

and sends the first message t to the verifier V .

2. Afterwards, the verifier V carries out the following steps:

The verifier V chooses a random value c , wherein $c \in_{\mathbb{U}} [0, e - 1]$, and sends the random value c to the prover P .

3. Afterwards, the prover P carries out the following steps:

The prover P computes the response value $res = r \cdot \sigma^c \pmod{N}$, and sends the response value res to the verifier V .

4. Finally, the verifier V carries out the following step:

The verifier V outputs the value of the predicate $res^e \stackrel{?}{=} t(H(m))^e \pmod{N}$

[0046] This protocol provides a security assurance of $1/e$ which is about 2^{-16} in the TCG case. If the protocol is run interactive, one should repeat it at least 4 times. If the protocol is run in non-interactive mode, i.e., when a hash function / Fiat-Shamir heuristic is applied to compute the challenge, one should run the protocol about 10 times.

Specification of a privacy friendly attestation protocol

[0047] It is denoted the set T of publicly known PCR values that are considered correct by $T \doteq \{PCR_1, \dots, PCR_n\}$. The prover's AIK signing- / verification - key pair is denoted by (S_{AIK}, V_{AIK}) . The signature σ on the message m is $\sigma \doteq \text{sign}(S_{AIK}, m)$, which can be verified by evaluating the predicate $\text{verify}(V_{AIK}, \sigma)$. More precisely, the verification key V_{AIK} consists of the integers (N, e) where N is a RSA modulus and e is a prime exponent.

[0048] Let $A_V \doteq \{i_1, \dots, i_m\}$ the set of indices of the PCR values for which the prover P knows a signature. That is, the prover P knows a set of signatures S :

$$S \doteq \{\sigma_{i_1} \doteq \text{sign}(S_{AIK}, PCR_{i_1}), \dots, \sigma_{i_m} \doteq \text{sign}(S_{AIK}, PCR_{i_m})\}$$

[0049] It is A_V an index set, and $A_1, \dots, A_k \subset \{1, \dots, n\}$ arbitrary index sets. It is noted that for the PCR values described by the index sets A_1, \dots, A_k the prover P does not need to know the corresponding signatures. Finally G denotes the set of index sets $\{A_V, A_1, \dots, A_k\}$.

[0050] The above described privacy friendly attestation protocol allows the prover P to assert to the verifier V that it knows signatures on the PCR values described by one of the index sets $\{A_V, A_1, \dots, A_k\}$ in the set G . Thereby the verifier V does not get any information for which of the index sets $\{A_V, A_1, \dots, A_k\}$ in the set G the prover P knows the signature or signatures σ .

[0051] In the following a protocol for implementing the invention is described. Common inputs to the prover P , and the verifier V are: $C \doteq \{0, \dots, c^*\}$, the set of public known correct PCR values T , the set G , and the verification key V_{AIK} .

Private inputs to the prover P are: the set of signatures S, and the index set A_V . The joint computation of the prover P, and the verifier V is as follows.

1. The prover P carries out the following steps:

a) The prover P chooses for each $j \in A_V$ a random value r_j wherein $r_j \in Z_N^*$, and sets the value t_j to:

$$t_j \doteq r_j^e \bmod N.$$

b) The prover P chooses for $j \in \bar{A}_V$ a response value $res_j \in_U Z_N^*$, $c \in_U C$, and sets the value t_j to:

$$t_j \doteq res_j^e (H(PCRj))^c \bmod N$$

It is noted that the index set \bar{A}_V is the complement of the index set A_V in $\{1, \dots, n\}$.

c) The prover P sends the messages (t_1, \dots, t_n) to the verifier V.

2. After that, the verifier V carries out the following steps:

The verifier V chooses a random value c wherein $c \in_U C$, and sends the random value c to the prover P.

3. After that, the prover P carries out the following steps:

$$\tilde{C} \doteq \text{complete}(c, \{c_i : i \in \bar{A}_V\}, \Gamma^*).$$

a) The prover P computes for $j \in A_V : c_j \in \tilde{C}$ the response value $res_j \doteq r_j \sigma_j^{c_j} \bmod N$.

b) The prover P sends $((res_1, c_1), \dots, (res_n, c_n))$ to the verifier V.

Remark: For $j \in \bar{A}_V$ the (res_j, c_j) were already computed in the round 1.

4. The verifier V carries out the steps:

The verifier V outputs the value of the following predicate:

$$\left(\bigwedge_{j=1, \dots, n} s_j^e \stackrel{?}{=} t_j (H(PCRj))^{c_j} \bmod N \right) \wedge \text{is-consistent}(c, \{c_1, \dots, c_n\}, \Gamma)$$

wherein \wedge denotes an AND-operation.

[0052] In the above described protocol the random value c is interpreted as secret and c_i as share of the secret c and the c_i are the challenges in the zero knowledge protocol with which is proofed that a signature on the PCR $_i$ value is known.

[0053] As mentioned in the previous section, $c^+ < e$ should hold for this proof to work. Also, the secret sharing scheme θ should be compatible with this. In practice, one would run many, e.g., 10 instances of this protocol in parallel and determine c via a hash function of all the first messages, i.e., the tuples (t_1, \dots, t_n) , the involved AIK and all the PCR values. More precisely, the one would concatenate for each i the c_i (about 16 bits each) from all the protocol instances, to obtain a larger \hat{C}_i (about 160 bits), then use the hash function to determine \hat{C} (e.g., 160 bits), use the secret sharing scheme to obtain the remaining \hat{C}_i , (e.g., 160 bits), split then into smaller parts (e.g., 16 bits) and use each of these small parts with on the parallel protocol instance.

[0054] Which configuration values are permissible configuration values is defined by the verifier V. For example, the verifier V can mark a configuration value as not permissible if the corresponding configuration does not fulfill the desired security properties.

[0055] Abuse might arise if a configuration value is indicated as permissible by the verifier V but does not represent a possible configuration. If the verifier V transmits such improper configuration or dummy values together with a proper configuration value PCRs actually representing a possible configuration to the prover P the verifier V might spy out the

actual configuration of the prover P. Because the set of transmitted configuration values comprises mainly dummy values, the prover P can only prove to the verifier V that its configuration is the configuration corresponding to the configuration value PCRs and with that the verifier V knows the actual configuration of the prover P. If the prover P cannot prove that its configuration is in the set of transmitted configuration values, i.e. its configuration does not correspond to the configuration value PCRs, the verifier V can transmit a further slightly modified set of configuration values to the prover P and repeat the procedure until it has found out the actual configuration of the prover P. To prevent this kind of abuse the method according to the invention can comprise the following step. After the prover P has received the permissible configuration values PCR1 ... PCRn the prover P checks whether configurations having the configuration values PCR1 ... PCRn actually may exist, and if this is the case the integrity proof is processed. Otherwise the prover P can for example abort the proof or ask the verifier V for a new set of configuration values.

[0056] A similar abuse by the verifier V may occur if the verifier V transmits only a very small number of configuration values to the prover P. Also in this case the verifier V has the possibility to spy out the actual configuration of the prover. If for example the verifier V transmits only one configuration value to the prover P, the prover P can only prove whether its configuration corresponds to transmitted configuration value or not. In the case the configuration of the prover P corresponds to the transmitted configuration value the prover P proves this to the verifier V and with that the verifier V has learned the configuration of the prover. To avoid this, the prover P can check whether the number of transmitted configuration values does not fall under a minimum number of transmitted configuration values. If the number of transmitted configuration values exceeds the necessary minimum number the integrity proof is processed. Otherwise the prover P can for example abort the proof or ask the verifier V for a bigger set of configuration values.

[0057] Computer readable program code or program code in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b) reproduction in a different material form.

Claims

1. Method for privacy-protecting integrity attestation of a computing platform having a trusted platform module (TPM), comprising:
 - a) receiving configuration values (PCR1 ... PCRn),
 - b) determining by means of the trusted platform module (TPM) a configuration value (PCRp) depending on the configuration of the computing platform (P),
 - c) signing by means of the trusted platform module (TPM) the configuration value (sign(PCRp)), and
 - d) in the event that the configuration value (PCRp) is one of the received configuration values (PCR1 ... PCRn), the computing platform (P) proving to a verifier (V) that it knows the signature (sign(PCRp)) on one of the received configuration values (PCR1 ... PCRn).
2. Method according to claim 1, wherein the proof is carried out by a cryptographic proof.
3. Method according to claim 1 or 2, wherein the proof is carried out by a zero-knowledge proof.
4. Method according to any of the previous claims 1 to 3,
 - wherein after receiving the configuration values (PCR1 ... PCRn) the computing platform (P) checks whether configurations having the configuration values (PCR1 ... PCRn) actually may exist, and
 - if this is the case the method is further processed.
5. Method according to any of the previous claims 1 to 4, wherein the computing platform (P) checks whether the number of received configuration values exceeds a minimum number of configuration values, and
 - if this is the case the method is further processed.
6. Computer program comprising program code for performing the steps of the method according to any of the previous claims 1 to 5 when loaded in a digital processor of a computer.

7. Computer program product stored on a computer usable medium, comprising computer readable program code for causing a computer to perform the steps of the method according to any of the previous claims 1 to 5.

5 8. A computing platform for privacy-protecting integrity attestation, the computing platform (P) being adapted to receive configuration values (PCR1 ... PCRn), comprising:

10 a trusted platform module (TPM) that is adapted to determine a configuration value (PCRp) in dependence on the configuration of the computing platform (P) and to sign the configuration value (sign(PCRp)), and wherein the computing platform (P) comprises the functionality that in the event that the configuration value (PCRp) is one of the received configuration values (PCR1 ... PCRn) the computing platform (P) is adapted to prove to a verifier (V) that it knows the signature (sign(PCRp)) on one of the received configuration values (PCR1 ... PCRn).

15 9. A network for privacy-protecting communication, comprising:

- a computing platform (P) having a trusted platform module (TPM), and
- a verifier (V) connected to the computing platform (P),

20 wherein the computing platform (P) is adapted to receive configuration values (PCR1 ... PCRn), wherein the trusted platform module (TPM) is adapted to determine a configuration value (PCRp) in dependence on the configuration of the computing platform (P) and to sign the configuration value (sign(PCRp)), and wherein the computing platform (P) comprises the functionality that in the event that the configuration value (PCRp) is one of the received configuration values (PCR1 ... PCRn) the computing platform (P) is adapted to prove to the verifier (V) that it knows the signature (sign(PCRp)) on one of the received configuration values (PCR1 ... PCRn).

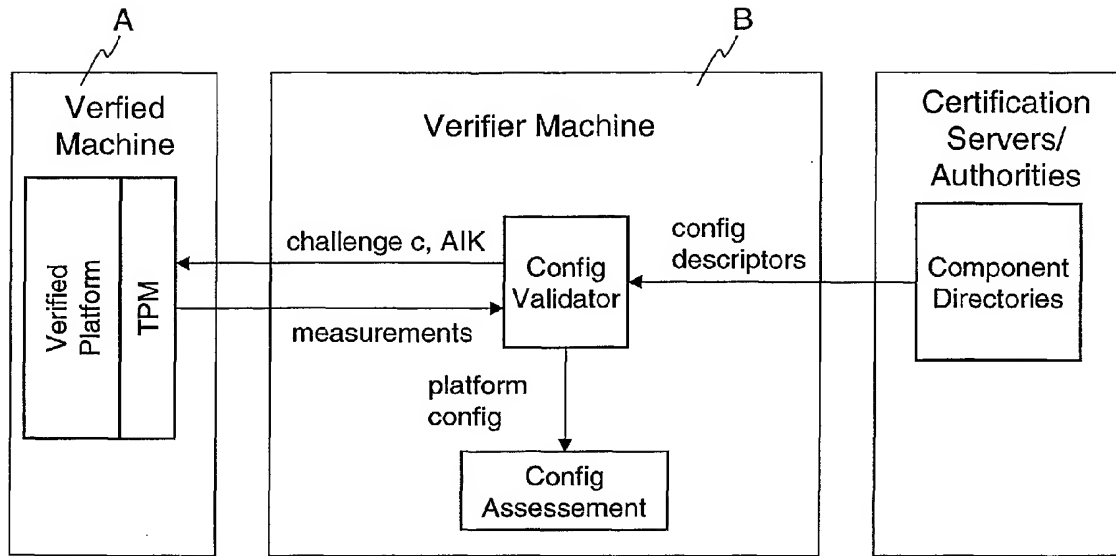


Fig. 1

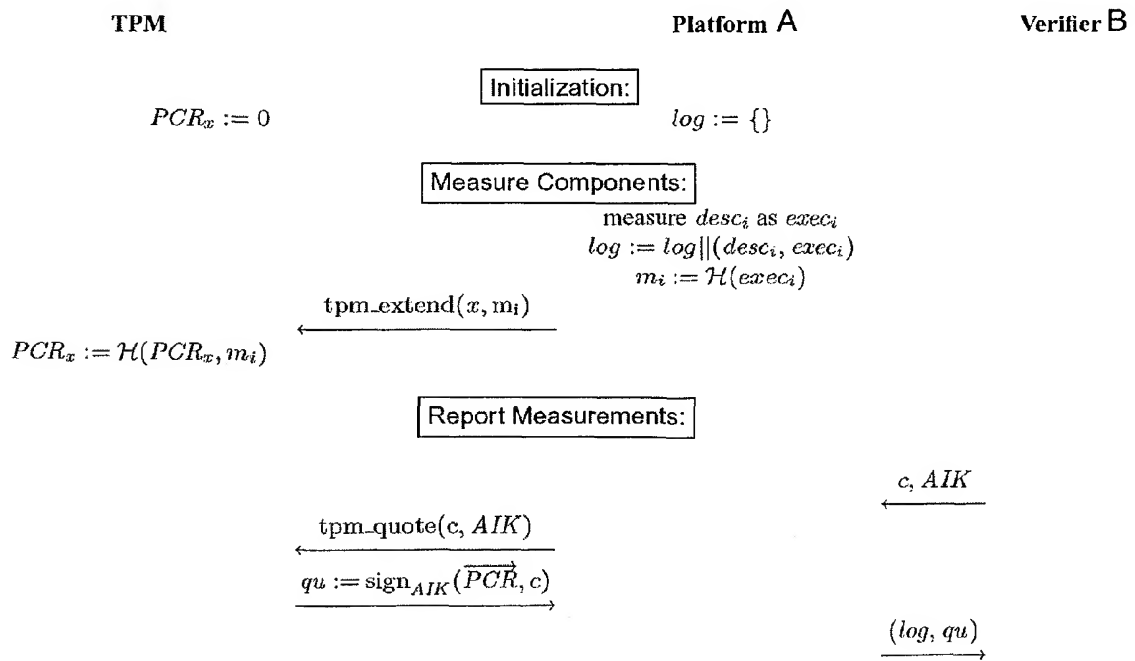


Fig. 2

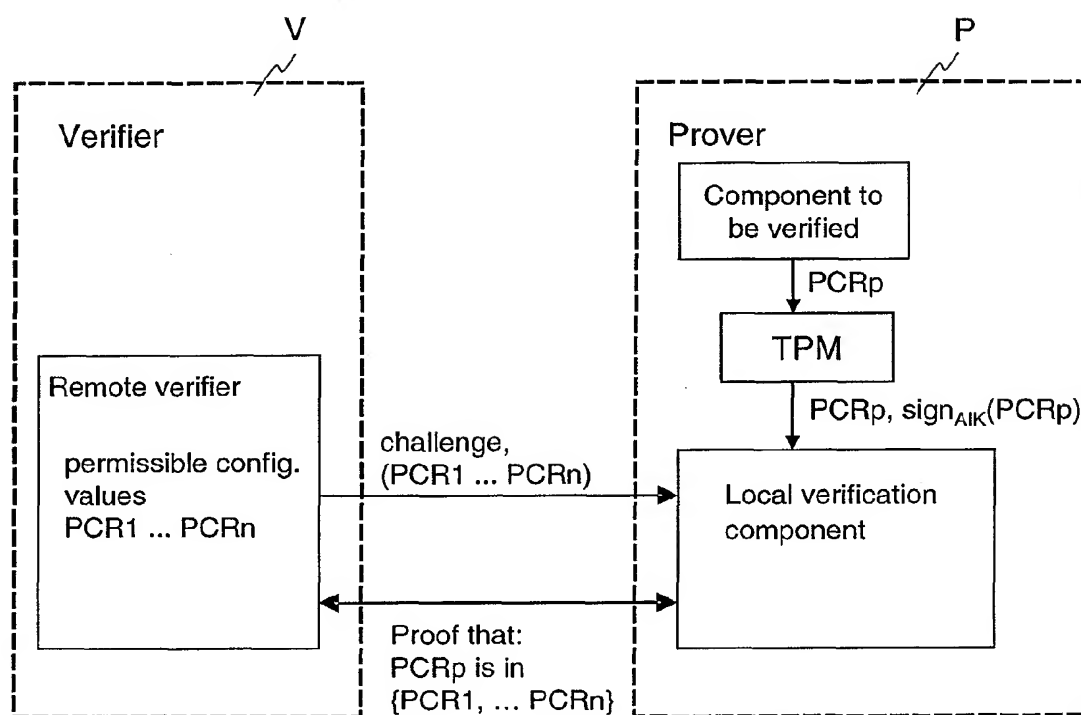


Fig. 3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 04 40 5442

| DOCUMENTS CONSIDERED TO BE RELEVANT | | | |
|--|---|---|--|
| Category | Citation of document with indication, where appropriate, of relevant passages | Relevant to claim | CLASSIFICATION OF THE APPLICATION (Int.Cl.7) |
| X | US 2004/064457 A1 (FISH ANDREW J ET AL) 1 April 2004 (2004-04-01) | 1,4-9 | H04L9/32 G06F21/00 |
| Y | * abstract * * figures 3-7 * * paragraph [0009] * * paragraph [0011] - paragraph [0013] * * paragraph [0023] * * paragraph [0026] - paragraph [0030] * * paragraph [0048] - paragraph [0049] * * paragraph [0064] - paragraph [0068] * | 2,3 | |
| Y | US 2004/103281 A1 (BRICKELL ERNIE F) 27 May 2004 (2004-05-27) * figure 5 * * paragraph [0003] * * paragraph [0010] * * paragraph [0015] * * paragraph [0018] - paragraph [0023] * * paragraph [0027] * * paragraph [0038] * * paragraph [0041] * | 2,3 | |
| A | US 2003/074548 A1 (CROMER DARYL CARVIS ET AL) 17 April 2003 (2003-04-17) * abstract * * figure 1 * * paragraph [0003] * * paragraph [0005] * * paragraph [0008] - paragraph [0019] * * paragraph [0025] * * claims 1,3,4,11,12,14,21,22,24 * | 1-9 | |
| A | US 5 633 929 A (KALISKI JR BURTON S) 27 May 1997 (1997-05-27) * column 5, line 58 - line 65 * * claims 1,2,4,6 * | 2,3 | |
| The present search report has been drawn up for all claims | | | |
| Place of search Munich | | Date of completion of the search 9 December 2004 | Examiner Kopp, K |
| CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document | | T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document | |

1
EPO FORM 1503 03.92 (P04001)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 04 40 5442

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

09-12-2004

| Patent document cited in search report | Publication date | Patent family member(s) | Publication date |
|---|---------------------|----------------------------|---------------------|
| US 2004064457 A1 | 01-04-2004 | NONE | |
| US 2004103281 A1 | 27-05-2004 | WO 2004051923 A1 | 17-06-2004 |
| US 2003074548 A1 | 17-04-2003 | NONE | |
| US 5633929 A | 27-05-1997 | NONE | |